REMARKS

Claims 1-19 stand rejected. Claim 1 has been amended to set forth the recited subject matter more clearly. Claim 2 has been amended to provide proper antecedent basis, in accordance with the amendment to claim 1. Reconsideration of the application in view of the remarks set forth below is respectfully requested.

Rejections Under 35 U.S.C. § 102

The Examiner rejected claims 1, 2, 9 and 10 under 35 U.S.C. §102(e) as being anticipated by Cho et al., (U.S. Patent No. 6,625,685), hereinafter Cho. The Examiner also rejected claims 1 and 2 under 35 U.S.C. §102(e) as being anticipated by Van Hook et al., (U.S. Patent No. 6,564,304) herein after Van Hook. With regard to independent claim 1, the Examiner stated:

> As for claims 1-2, Cho teaches a memory system and method of processing a read request comprising the acts of: transmitting a first read request from a requesting device, the first read request comprising a first system address **[transactions from the one of the devices (column 1, lines 12-13)]**; receiving the first read request at a memory controller **[memory controller receiving an access request}**; mapping the first system address into a first memory address, the first memory address comprising a first chip select, first bank address, a first row address and a first column address **[first portion of the address is a row address, the second portion of the address is a column address (column 1, lines 19-21); memory controller may provide a separate select signal (column 1, lines 26-29); portion of the address is used to select a bank (column 2, lines 44-45);** sorting the first read request by one of the first chip select and the first bank address such that the first read request in injected into a first read queue **[transaction queue 30 issues the transactions to one of the channel control circuits 32A-32B (column 7, lines 6-8; also see Fig. 2)].**
>
> As for claims 1-2, Van teaches a memory system and method of processing a read request comprising the acts of: transmitting a first and second read requests from a requesting

device, the first and second read requests comprising a first and second system addresses **[request from one of the memory master M0-M4 (see Fig. 1)]**; receiving the first and second read requests at a memory controller **[memory controller 105 (column 4, lines 65)]**; mapping the first and second system addresses into a first and second memory addresses, the first and second memory addresses comprising the first and second chip selects, first and second bank addresses, first and second row addresses and first and second column addresses **[Inherent in the art, because a bank address in addition to a row and column addresses is needed to specify a memory location; for Inherency support, please see patent numbers 6,470,733 and 5,367,669]**; sorting the first and second read requests by one of the first and second chip selects and the first and second bank addresses such that the first and second read requests in injected into a first and second read queues **[the requests from the RQ0-RQ4 are placed into sort queue 101 (odd or even 102 and 103)]**.

Applicants respectfully traverse these rejections. Anticipation under Section 102 can be found only if a single reference shows exactly what is claimed. *Titanium Metals Corp. v. Banner*, 778 F.2d 775, 227 U.S.P.Q. 773 (Fed. Cir. 1985). For a prior art reference to anticipate under Section 102, every element of the claimed invention must be identically shown in a single reference. *In re Bond*, 910 F.2d 831, 15 U.S.P.Q.2d 1566 (Fed. Cir. 1990). To maintain a proper rejection under Section 102, a single reference must teach each and every element or step of the rejected claim. *Atlas Powder v. E.I. du Pont*, 750 F.2d 1569 (Fed. Cir. 1984). Thus, if the claims recite even one element not found in the cited reference, the reference does not anticipate the claimed invention.

As disclosed in the present specification, when a memory request is sent, traditional memory controllers typically open a single bank and access a number of columns before closing the bank. Pg. 20, lines 13-18. Disadvantageously, running multiple successive accesses to row and columns in a single bank slows system performance due to the conflicts induced by repeated activity along the same buses and the continual targeting of the same

memory device. Pg. 20, lines 19-21. To facilitate increased memory bandwidth and reduced memory latency in accordance with one exemplary embodiment of the present invention, either the chip select address or the bank address may be used to sort each access request such that simultaneous processing is possible. Pg. 25, lines 8-10. By sorting requests into one of two queues, based on either the chip select or the bank address, the requests may be processed such that consecutive request processing is implemented by alternating between the two queues such that consecutive cycles are processed to different banks or different memory devices (as determined by the chip select). Pg. 24, lines 15-16; Pg. 25, lines 18-20; Pg. 27, lines 9-14.

Accordingly, claim 1, as amended, recites "sorting the first read request by one of the first chip select and the first bank address such that the first read request is injected into a selected one of a first read queue and a second read queue based on the sorting by the one of the first ship select and the first bank address." Neither Cho nor Van Hook discloses this element.

Contrary to the Examiner's assertion, the Cho reference *does not* disclose sorting a request into a particular queue based on the bank address or chip select of the request. The Examiner directs Applicants to Col. 7, lines 6-8 and Fig. 2 of Cho. As explicitly stated in the cited passage, the "transaction queue 30 is configured to receive and queue memory transactions from bus 24, and to issue those transactions to one of channel control circuits 32A-32B." Col. 7, lines 6-8. As disclosed by Cho and clearly illustrated in Fig. 2, the requests are *not* sorted into one of two queues. The requests are simply received at a single transaction queue 30. Because Cho only discloses a single transaction queue 30, it is clear that Cho does not disclose sorting requests between a first queue and a second queue, much

less that the sorting and selection of the queue into which the request is injected is based on either the chip select or the bank address of the address.

If the Examiner is asserting that the channel control circuits 32A and 32B disclosed in Cho are analogous to the presently recited first and second read queues, Applicants note that the channel control circuits 32A-32B are not queues but are configured to receive transactions from the transaction queue 30. The channel control circuits 32A-32B are simply configured to provide access to a corresponding channel 34A-34B in response to the memory transaction provided from the transaction queue 30. Col. 6, lines 18-34. Each channel control circuit 32A-32B includes configuration registers 36A-36B which are used to program the configuration of the memory system. Col. 6, lines 34-35. At best, Cho discloses sorting (received from a single queue) between two channels, *not* sorting requests between queues. Further, the Cho reference does not disclose sorting a request by a chip select or bank address, as further recited in claim 1. Accordingly, the Cho reference cannot possibly anticipate the subject matter recited in claim 1.

Likewise, Van Hook does not disclose sorting requests into a first queue or a second queue, depending on the chip select or bank address, as recited in claim 1. As with Cho, Van Hook simply discloses delivering all requests to a single sort queue 101. Col. 5, lines 9-10. The sort queue 101 is implemented to provide temporary storage such that the requests in the sorting queue can be reordered. *See* e.g., Fig. 2; Col. 5, lines 44-53. The reordering involves grouping reads together with reads and writes together with writes and grouping multiple requests together with other requests to a single bank. Col. 5, lines 28-30, lines 49-51. Accordingly, Van Horn *does not* disclose sorting between a first and second queue based on

the bank address or chip select, as recited in claim 1. Accordingly, Van Horn cannot possibly anticipate the recited subject matter.

For at least the reasons set forth above, it is clear that neither Cho nor Van Horn discloses sorting requests between a first and second read queue, much less sorting the requests by chip select or bank address. Specifically, because neither of the references disclose "sorting the first read request by one of the first chip select and the first bank address such that the first read request is injected into a selected one of a first read queue and a second read queue based on the sorting by the one of the first ship select and the first bank address," as recited in claim 1, neither of the references can possibly anticipate the subject matter recited in claim 1. Accordingly, Applicants respectfully request withdrawal of the Examiner's rejection and allowance of claims 1, 2, 9 and 10.

Rejections Under 35 U.S.C. § 103

The Examiner rejected claims 2-8, 11-19 35 U.S.C. § 103(a) as being unpatentable over Cho or Van in view of Tran et al. (U.S. Patent No. 6,598,132), hereinafter Tran. Specifically, the Examiner stated:

> As for claims 3-8, and 11-19, Cho and Van disclose the claimed invention, including that the second chip select identifies a corresponding DIMM to which the second request is directed and a second bank address identifies a corresponding bank in a corresponding DIMM to which the second request is directed **[by asserting the chip select, the corresponding DIMM is selected (column 7, lines 47-49); selecting a bank within the SDRAM [see column 7, lines 38-40],** however Cho or Van fails to specifically teach alternating the read requests from the first and second queues, so that back-to-back requests are processed between the first and second queues and back-to-back request are not processed to the same bank.
>
> Tran discloses a buffer memory, a buffer manager and a queue manager, Tran further discloses a SDRAM including

> number of banks and an order in which the buffer manager read and writes to the various banks. Furthermore Tran discloses that back-to-back read are not processed for the same bank, since each bank requires additional cycles to recover after a read access, so that it would be advantageous to not process back-to-back read requests to the same bank in order to have time to precharge the bank for the next access **[column 6, lines 34-65].**
>
> Accordingly, it would have been obvious to one having ordinary skill in the art at the time of the current invention to use bank accesses as being taught by Tran into Cho and Van memory system in order to have time to precharge the bank for the next access.

Applicants respectfully traverse this rejection. The burden of establishing a *prima facie* case of obviousness falls on the Examiner. *Ex parte Wolters and Kuypers*, 214 U.S.P.Q. 735 (PTO Bd. App. 1979). Obviousness cannot be established by combining or modifying the teachings of the prior art to produce the claimed invention absent some teaching or suggestion supporting the combination or modification. *See ACS Hospital Systems, Inc. v. Montefiore Hospital*, 732 F.2d 1572, 1577, 221 U.S.P.Q. 929, 933 (Fed. Cir. 1984). Accordingly, to establish a *prima facie* case, the Examiner must not only show that the combination includes *all* of the claimed elements, but also a convincing line of reason as to why one of ordinary skill in the art would have found the claimed invention to have been obvious in light of the teachings of the references. *Ex parte Clapp*, 227 U.S.P.Q. 972 (B.P.A.I. 1985).

Claim 3, which is dependent on claim 1, recites "alternately selecting read requests from the first read queue and the second read queue such that back-to-back requests are alternately processed between the first read queue and the second read queue." Similarly, independent claim 11 recites "prioritizing the processing of the plurality of read requests such that back-to-back read requests are not directed to the same one of a corresponding chip select and a bank address." Independent claim 17 recites "an arbiter operably coupled to each

of the first and second read queues and configured to select requests stored in the queues such that consecutive read requests do not have the same one of the chip select in the bank address."

As recognized by the Examiner, neither Cho nor Van disclose sorting read requests between first and second queues such that back-to-back requests are not processed to the same bank. However, contrary to the Examiner's assertion, the Tran reference *does not* cure the deficiencies of the Cho and Van references, because it does not disclose this feature. The Examiner directed Applicants to Fig. 3 of Tran. However, Fig. 3 does not even disclose a first read queue and a second read queue or sorting requests into separate read queues, much less that sorting between queues is done based on chip select or bank address such that requests can be alternately processed from the queues.

Further, while the passage cited by the Examiner may indeed disclose that memory banks require a certain amount of time to recover (precharge) before any address of that bank can be read or write accessed again, (Col. 6, lines 51-55), as also discussed by Applicants in the present specification, Tran *does not* disclose separately accessing alternating queues wherein requests have been sorted into the queues by chip select or bank address. In other words, while the Tran reference clearly recognizes that there are delays associated with processing back-to-back requests to the same bank, the Tran reference *does not* disclose the solution disclosed by Applicants and recited by the present claims. Thus, contrary to the Examiner's assertion, the Tran reference does not disclose sorting read requests into one of a first or second read queue based on the bank address or chip select much less that the sorting is done such that read requests are processed alternately between the queues such that

consecutively processed requests do not have the same chip select and bank address, as recited in claims 3, 11 and 17.

Because none of the references, either alone or in combination discloses each of the elements recited in the present claims, the present claims cannot possibly be rendered obvious by the cited combination. Accordingly, Applicants respectfully request withdrawal of the rejection and allowance of claims 3-8 and 11-19.

Conclusion

In view of the remarks and amendments set forth above, Applicants respectfully request allowance of claims 1-19. If the Examiner believes that a telephonic interview will help speed this application toward issuance, the Examiner is invited to contact the undersigned at the telephone number listed below.

Respectfully submitted,

Date: May 13, 2004

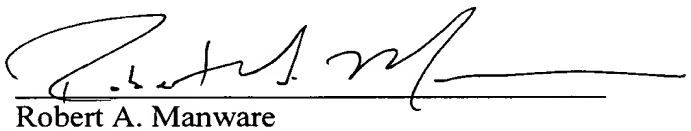

Robert A. Manware
Reg. No. 48,758
(281) 970-4545

Correspondence Address:
Hewlett-Packard Company
IP Administration
Legal Department, M/S 35
P.O. Box 272400
Fort Collins, CO 80527-2400